arXiv:1605.02240v1 [cs.CV] 7 May 2016

On Image segmentation using Fractional Gradients- Learning Model Parameters using Approximate Marginal Inference

Anish Acharya
Schlumberger
AAcharya2@slb.com

Dr. Uddipan Mukherjee
Intel Corporation
umukherjee@uci.edu

Dr. Charless Fowlkes
University of California Irvine
fowlkes@uci.edu

Abstract

Estimates of image gradients play a ubiquitous role in image segmentation and classification problems since gradients directly relate to the boundaries or the edges of a scene. This paper proposes an unified approach to gradient estimation based on fractional calculus that is computationally cheap and readily applicable to any existing algorithm that relies on image gradients. We show experiments on edge detection and image segmentation on the Stanford Backgrounds Dataset where these improved local gradients outperforms state of the art, achieving a performance of 79.2% average accuracy.

Keywords

Image Segmentation, Fractional Derivative

1. Introduction

A huge variety of modern computer vision algorithms are built on the computation of local derivatives of image brightness which provide local discriminative cues to scene structure, e.g., texture based segmentation [20], HOG based [8], and SIFT [19] frameworks. Typically an input image is convolved with an image filter bank consisting of gradient filters in different orientations and scales. The response of these filters are used to construct a feature space which embeds structural image information. This approach is found to be incredibly powerful for detecting key image features or recognizing objects over different scales, appearances and poses. The success of these algorithms prove the importance of gradient orientations as a robust feature for recognition [28, 21, 22, 3]. To date the most widely used gradient based features are extracted using the First Order Derivative of a Gaussian (DOG) kernel taken in different orientation and scales or considering the zero crossing in the response to the Laplacian of Gaussian Operator (LOG). As the order of the gradient operator increases, the response to edges is also higher, but an increased order also gives an undesirably higher response to noise or isolated points. In order to avoid noise enhancement images are often smoothed prior to edge detection. However, smoothing can cause a loss of structural information in some cases. Thus, choosing the proper gradient order for a given image depends largely on the image, e.g. we might want a weaker edge response for a noisy image and a stonger one for a crisp image.

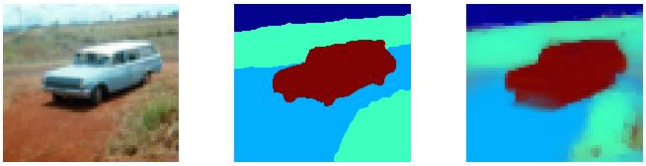

Figure 1. Image segmentation using FDOG edge detector. From left to right: original image, true label and estimated label.

In order to address this problem we propose a framework for generating a family of gradient operators with fractional orders, and show that the fractional gradients work better than integral ones for edge detection and eventual image segmentation. In order to design the fractional gradient operators, we use concepts from Fractional Calculus, which has been widely and efficiently used in the fields of signal processing. The fractional orders introduce higher order hyper-damped and ultra-damped poles which lends much more flexibility and accuracy in modeling complicated system dynamics [1, 9].

In essence, we use fractional order gradients for designing fractional order derivative of Gaussian (FDOG) filters and show that our design achieves better edge detection performance than existing integer order gradient filters (Figure 2). In order to validate this performance gain, we apply our method to image segmentation (on images from Stanford Backgrounds Dataset [15], an example shown in Figure 1) by building Histogram of Gradient Orientation (HOG) features on top of the FDOG filters, and learning a Markov Random Field based Model using Tree Reweighted Belief Propagation based on the approach proposed in [10].

The rest of the paper is organised as follows: Section 2

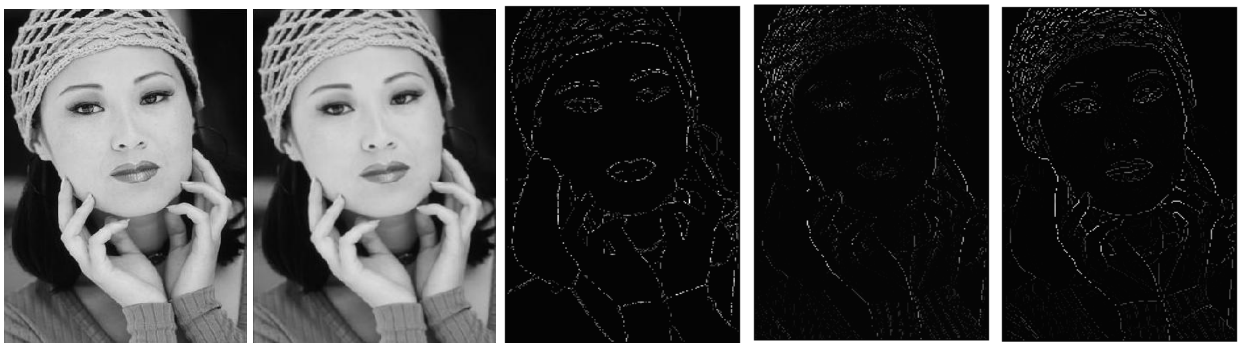

Figure 2. Edge detection using FDOG: Left to right: sample image, smoothed image, manual edge detection, DOG edge detection and FDOG with v=0.5. It can be observed that FDOG performance is closer to manual detection than DOG.

gives a detailed description of the design of FDOG filters and demonstartes its performance on test images. Section 3 describes how the FDOG based approach can be used for image segmentation and Section 3.1 depicts the performance gain for the same over existing methods, followed by a conclusion in Section 4.

2. Fractional Order Derivative of Gaussian Filters for Edge Detection

The most fundamental structural information embedded in an image is its edges. Edge detection is a well established notion in Computer Vision and is still a significant area of research. Among the various algorithms used for edge detection, detectors like Canny [5], Sobel, Pretwitt, Robert's Cross are particularly popular, Canny being the most reliable and widely used in the community. Recently, the works of Perona, Malik et.al., based on Anisotropic Diffusion algorithm [26], and a new edge detection scheme proposed by Harris et.al. [13], also gained popularity. However, one major disadvantage of Harris Detector is that it is not scale invariant. Though there has been many recent advances on the improvement of this [4, 7], most of these techniques are computationally complex compared to the Canny as well as Anisotropic Diffusion based detectors. These detectors might be extremely efficient when the sole task is edge detection. However, when we are concerned with a broader problem such as segmentation, edge detection is one of the many steps involved, and we need a robust and efficient detector without much overhead of computational complexity.

An optimal edge detector captures the underlying geometric structure of the image (i.e. detects all the edges successfully) without being much affected by the noise present. As has been already shown by Canny et.al. [5], it is impossible to realize such an optimal filter for edge detection, although the first order gradient of Gaussians is a good approximation with around 20% error. The order of the filter gradient is directly proportional to the number of detected edges and the amount of detected noise. Thus we get more peaked response at an edge with the second order LOG (Laplacian of Gaussian) as compared to the first order DOG (derivative of gaussian), with an additional side-effect of a higher response to noisy peaks for LOG as compared to DOG. In order to control this trade-off more efficiently, we propose a detector with fractional order gradient operator (FDOG) and show that the edge detection performance can be improved by using FDOG.

The FDOG filters are designed using concepts from fractional order calculus. Similar filters have been designed in the frequency domain and has been practiced frequently in various control applications. Acharya et.al. [1] proposed a structured exact implementation of arbitrarily chosen cut-off based filters. We extend this idea to design a 2-D convolution mask based on the general discrete approximation of the n-th order differentiation operator, represented by (1).

$$\begin{aligned}\frac{\partial^v s(x,y)}{\partial x^v} &= s(x,y) - vs(x-1,y) + \frac{v(v-1)}{2}s(x-2,y) \\ &\quad - \frac{v(v-1)(v-2)}{6}s(x-3,y) + ..., \\ \frac{\partial^v s(x,y)}{\partial y^v} &= s(x,y) - vs(x,y-1) + \frac{v(v-1)}{2}s(x,y-2) \\ &\quad - \frac{v(v-1)(v-2)}{6}s(x,y-3) + ... \end{aligned} \quad (1)$$

The gradient to be computed is thus given by the sum of the partial derivatives (Equation 2).

$$\bigtriangledown^v = \frac{\partial^v s(x,y)}{\partial x^v} + \frac{\partial^v s(x,y)}{\partial y^v} = \bigtriangledown_x^v + \bigtriangledown_y^v \quad (2)$$

where instead of restricting the order to be integer i.e. in case of DOG, $v = 1$ and LOG $v = 2$; we propose to let v take any fractional value. Theoretically, v should be limited to one decimal place to avoid mathematical intractability. Also, for efficiency, we restrict the number of terms in equation 1 to three. Figure 2 shows the result of applying FDOG of order 0.5 for edge detection on a standard

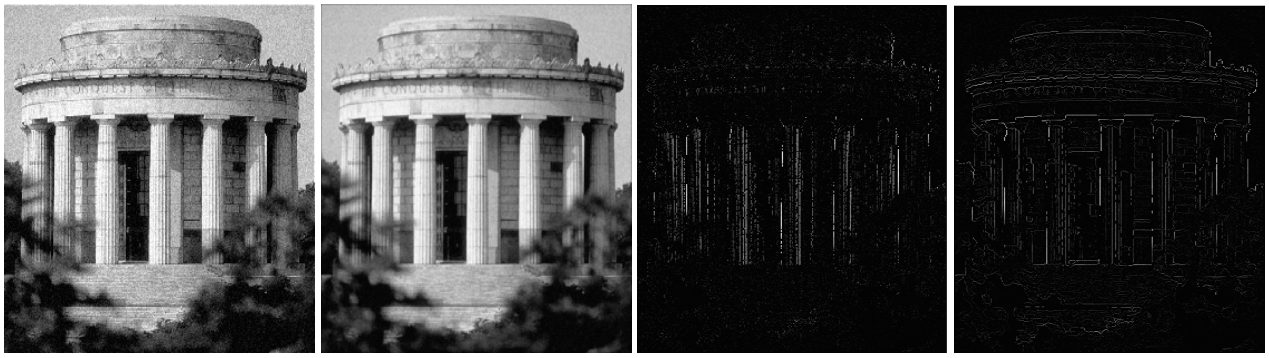

Figure 3. Edge detection using FDOG on a noisy image. From left to right: original image, smoothed with $\sigma = 2$ Gaussian kernel, DOG edge detection and FDOG edge detection. It can be observed that FDOG outperforms DOG even on a noisy image.

image taken from the Berkeley Segmentation Dataset. Figure 2 also shows a DOG kernel based edge detection and a benchmark obtained by averaging over the five hand edge-detected versions of the same image. All the results are subjected to non maximal suppression and displayed in normalized scale for comparison purposes. As can be observed, FDOG of order 0.5 seems to perform better than standard DOG operator visually. In other words, FDOG weeds out noise and reveals structural image information more clearly. Even in highly noisy images, FDOG seems to perform better, as is verified on an image (taken from the CSIQ image database [18]) corrupted with white noise (Figure 3).

2.1. Evaluating FDOG Edge detection

The standard measure to evaluate the performance of an edge detector is to measure the Peak Signal-to-noise Ratio (PSNR) [5]. This is an effective measure when the noise models of an image are provided, or more generally if the imaging circumstances are known before hand and the detected edges can be clearly classified as true edges or noise. A more practical and widely adopted approach for any arbitrary given image is to filter the image and measure the structural loss caused therby. It assumes a general framework of the corrupted image as given in equation 3.

$$x[n] = s[n] + w[n] \tag{3}$$

where $x[n]$ is the given image and $s[n]$ is the filtered one. The PSNR is then defined as shown in Equation 4.

$$\begin{aligned} PSNR &= 10log_{10}\frac{max(s^2[n])}{MSE}, \\ MSE &= \frac{1}{MN}\sum_{i=1}^{M}\sum_{j=1}^{N}[\hat{s}(i,j) - s(i,j)]^2 \end{aligned} \tag{4}$$

where $s[n]$ is the maximum signal value, typically 255 when considering image pixel values, M and N are the height and width of the image, $\hat{s}$ is the processed image and s is the original image, and MSE is the mean squared error.

While PSNR gives us an idea of how much structural information is retained in the gradient operated image, another criterion of good detection is to minimize the effect of noise. This can be addressed using the idea of False Discovery rate (DE) [2], which is given in equation 5.

$$DE = P_m * P_f \tag{5}$$

where P_m and P_f are respectively Probability of missed detections and probability of false positives. A low value of this product implies good detection i.e. better edge detector.

2.1.1 Finding the optimal Gradient order

The correct order of the gradient to be used in FDOG is data specific and should be optimized over all the images in the dataset. Thus finding the optimal order of the derivative can be now viewed as an optimization task with th score function, $J = \frac{PSNR}{DE}$ to be maximized. Any standard approach like gradient descent can be used to find the optimal order. We restrict the search space to within one decimal place i.e. v can only take values in steps of 0.1 and select the order with highest score. Applying our algorithm to various datasets like the berkley segmentation dataset [23], The Stanford backgrounds dataset [11] and CSIQ noisy database it is found that in most cases the optimal order turns out be in the range 0.5-0.8 with lower order for highly noisy images like astronomical images.

Figure 4 shows the plot of score function with respect to different filter orders applied on the Stanford background dataset. The score function for each order have been computed over all the images and averaged to produce the reported value. It can be seen that the approximate optimal order is somewhere around 0.5-0.8.

In order to validate our scoring scheme, we also generate a precision-recall curve for detection (Figure 5, Table 1). First we threshold the generated boundary map into several levels and at each level compute the precision, which is the probability that a machine generated pixel is a boundary

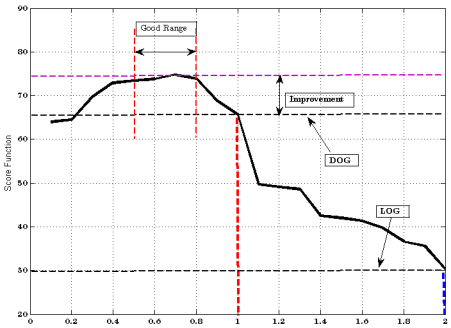

Figure 4. Performance of different Gradient Orders on Stanford Backgrounds Dataset

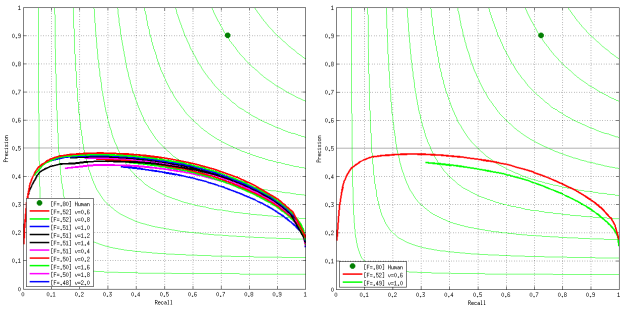

Figure 5. Top:left: Precision Recall curves on training BSDS500 for Varying Gradient Order right: comparing PR curves for FDOG order 0.6 and DOG

Gradient Order	ODS	OIS	AP
0.2	0.501	0.535	0.288
0.4	0.506	0.538	0.311
0.6	0.523	0.555	0.311
0.8	0.518	0.552	0.404
1	0.514	0.549	0.370
1.2	0.510	0.544	0.338
1.4	0.508	0.546	0.266
1.6	0.497	0.530	0.419
1.8	0.495	0.536	0.329
2	0.480	0.513	0.232

Table 1. Precision-recall and F-measures

pixel, and recall, which is the probability that a true boundary pixel is detected [23]. The precision-recall curve reflects the inherent trade-off between misses and false positives. It would be nice to be able to compare the curves for different algorithms in terms of a single number. In an attempt to do so we introduce the F-measure, which is the harmonic mean of precision and recall, as a yardstick for an algorithm's performance. In general, for two algorithms whose curves are roughly parallel, the one furthest from the origin dominates the other. The F-measure is defined at all points on the precision-recall curve. We report the maximum F-measure value across an algorithm's precision-recall curve as its summary statistic. These results should not be considered as edge detection results since these are based on just computing the raw local gradients with different orders using 1 and does not apply any processing before feeding to the benchmarking and thus the results would be lower than that in the BSDS leader-board. However, we use these PR curves as a way of comparing the descriptive power of different gradient orders to validate our formulated evaluation metric based on PSNR and FDR.

Figure 5 shows the precision-recall curves for varying orders and the variation of F-number with varying filter orders. Intuitively, and also as the curves show, an increased order enhances more noisy peaks and too small an order fails to mine structural information. From the training images it is observed that 0.6 gives the best F-score of 0.523 which is almost 1% improve over DOG. Whereas the LOG performs way bad and gives an F-score of 0.480 - almost 4% less than the best. We use the Filter order as 0.6 on the test images to generate the precision-recall curve and also provide the precision-recall of the standard DOG for comparison purposes in Figure 5.

Figure 6 shows the results of applying our FDOG filter with varying orders for edge detection.

3. Image Segmentation using FDOG

Image Segmentation is one of the most important application areas in Computer vision. Most of the image segmentation algorithms often use gradient based cues as features. In this section we apply the proposed fractional gradient based approach to an image segmentation task and evaluate its performance against state-of-the-art.

It is a common and successful practice to use Conditional Random Field (CRF) based approach in solving the segmentation task -it is also intuitively natural to think an image as a graph structure where each pixel can be considered as the nodes of the Graphical Model. CRFs are a probabilistic framework for labeling and segmenting structured data (equation 6), such as sequences, trees and lattices.

$$P(X \mid Y) = \frac{1}{Z(Y)} \prod_{C} \psi(X_c, Y) \prod_{i} \psi(X_i, Y) \qquad (6)$$

The underlying idea is that of defining a conditional probability distribution over label sequences given a particular observation sequence, rather than a joint distribution over both label and observation sequences.

However like most of the interesting problems exact inference is usually intractable or NP- hard problem due to the high tree-width of the graph structure. Learning the graph structure is also intractable in most cases [15]. Many previous research work has been focused on approximating the likelihood. We particularly focus on the recent work of Domke et.al. [10] where the Learning/ parameter fitting is

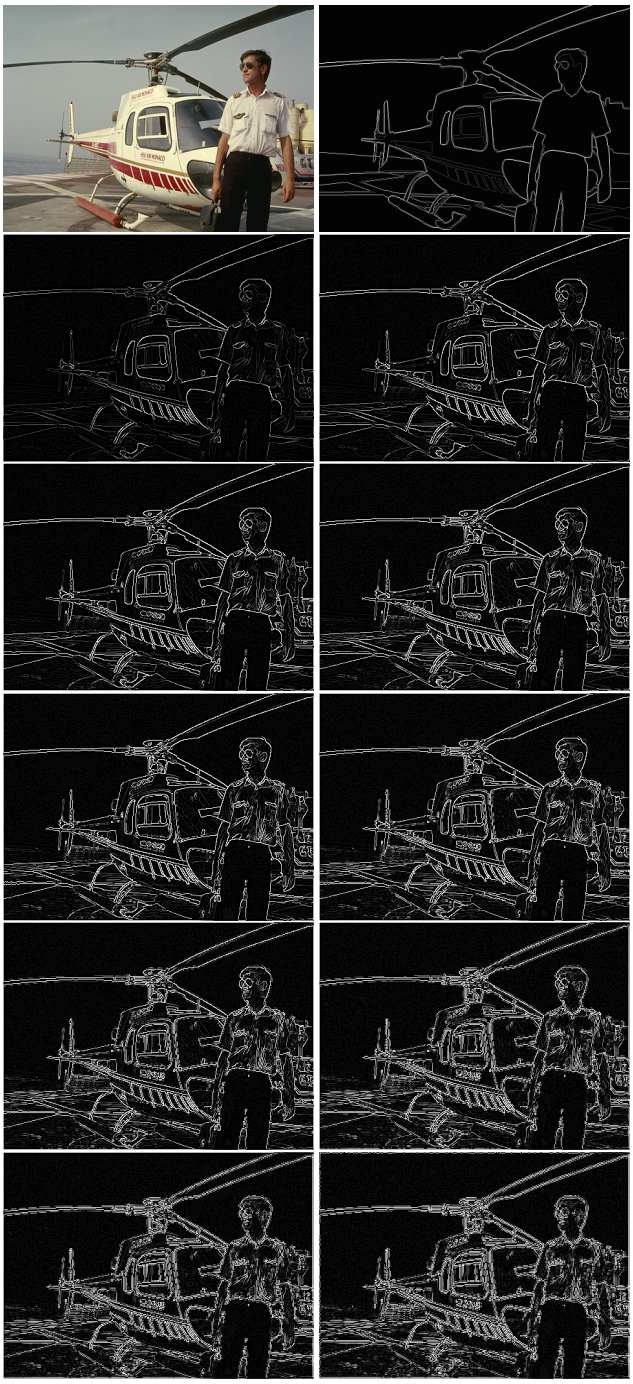

Figure 6. The Result of varying the Gradient Order. Top Row: Original image (left) and ground truth (right). Second row: FDOG with v=0.2 (left) and v=0.4 (right). Each row shows an increase in order in steps of 0.2 from left to right. Thus, the last row has FDOG wth v=1.8 (left) and v=2.0 (right).

viewed as repeated inference problem using the marginalization based loss function, which directly relates to the quality of prediction of a given marginal inference algorithm [10] also it seems to work better intuitively because it takes into account the approximation error and also it is robust to model miss-specification.

After learning a CRF model $P(Y \mid X)$ the query is to find the most likely configuration $\hat{x}$ given the observations. One option is to use the idea of Bayes estimator using the notion of utility function [25] which specifies the happiness of predicting $\hat{x}$ if X^* was the true output.Then the problem reduces to an optimization task where one finds the $\hat{X}$ that maximizes the Utility function.

$$\hat{X} = \underset{X}{\operatorname{argmax}} \sum_{x^*} P(X^* \mid Y) U(X, X^*) \tag{7}$$

It is also shown that using the utility function as the indicator function that is 1 when output equals X^* and zero otherwise, then it reduces to MAP estimate [25] i.e. if, $U(X, X^*) = I[X = X^*]$

$$\hat{x} = \underset{X}{\operatorname{argmax}}\, P(X \mid Y) \tag{8}$$

However, in real examples it might require humongous amount of data for accurate prediction. Thus, in practice, a utility function that maximizes the Hamming distance i.e. the number of components in the output that are correct, is used. This framework is popularly known as Maximum Posterior Marginal(MPM)[48]. In the context of the current problem of learning the graphical model parameters, we view Learning from an Empirical risk Minimization framework as proposed in [10]. where as usual the Risk is given by equation 9.

$$R(\theta) = \sum_{\hat{x}} L(\theta, \hat{x}) \tag{9}$$

where the Loss function $L(\theta, \hat{x})$ explains how good the fit is with parameter θ.

One can find exact marginals using tree structured graphs- which is not practical in case of graphs with high tree width such as the current problem in consideration because in that case finding the exact likelihood and its gradient becomes computationaly hard due to the log partition function and the marginals. One widely used approach to tackle this problem has been the use of markov chain Monte carlo to approximate the marginals or using Constructive Divergence based approaches[6][29] which also suffers from large complexity A more recent approach is to use approximate inference methods which can be viewed as approximating the partition function in the likelihood itself. Then it can be shown that the approximate marginals itself become the exact gradient of the Surrogate Loss. In recent literature Surrogate Loss seems to be widely used where the marginals being approximated by either Mean Field, Loopy Belief Propagation(LBP) or Tree-reweighted Belief Propagation (TRW).[31][33] If approximate log partition function is used that bounds the

true log partition function then Surrogate Likelihood is proved to be bounding the true likelihood where Mean field based Surrogate Likelihood Upper bounds and TRW based Surrogate Likelihood lower bounds. Other than these there is Expectation maximization(EM) based approaches which is appropriate in case of incomplete data[27], saddle point based approximations[30], pseudo-likelihood based approaches[14] and Piecewise Likelihood[17] etc. Another issue to be taken into consideration is model misspecification which often might be unavoidable due to the complex nature of problems or deliberately in case when the true model has many parameters to fit the data which leads to many degrees of freedom thus reduction becomes necessary.

For this particular problem we chose to train the model using a Clique Loss as in the present setting it is intuitively more straightforward to think about clique based loss rather than univariate loss as proved by Wainwright and Jordan et.al.[32] that due to the standard moment matching criterion of exponential families if the clique marginals are correct the joint distribution has to be correct though the joint distribution might be far off perfect even if the univariate marginals are perfect.It is defined mathematically as-

$$L(\theta, x) = -\sum_{C} log\mu(x_c; \theta) \tag{10}$$

which can be viewed minimizing the empirical risk of the mean KL-divergence of the true clique marginals to the predicted ones.

The next step is to do the inference which turns out to be tricky for graphs with high tree width which requires two operations- firstly the evaluation of the loss function, which is pretty straight forward and can be done simply by plugging the marginal obtained by running the inference algorithm into the loss function and secondly, one requires to know the gradient $\frac{dL}{d\theta}$. The gradient can be computed by solving a set of sparse linear equation which assumes that the optimization problem is exactly solved ,though in practice one has to set a threshold to truncate which can be done very elegantly with a much lower computational expense by forming the learning objective in terms of the approximate marginals obtained after a fixed number of iterations. [10]

It leads to a series of simple structured steps viz. Inputting parameters, applying the iterations of either TRW or mean field, computing predicted marginals, and finally the loss are all differentiable operations. Here we use back-propagating Tree Re-weighted Belief Propagation Framework which calculates the marginals and plug them into the marginal based loss after a certain number of iterations even if the inference algorithm did not converge, which is possible because each step is differentiable.

After execution of Back TRW the parameter $\theta(x_c)$ for the clique based loss can be represented mathematically as

$$\overleftarrow{\theta}(x_c) = \frac{dL}{d\theta(x_c)} \tag{11}$$

3.1. Experimental Results of using FDOG for segmentation

In order to experiment with our proposed FDOG based method we have chosen the Stanford background dataset which has 715 images of outdoor scenes taken from different publicly available datasets like LabelMe, MSRC, PASCAL VOC and Geometric Context; each image of resolution approximately 240 * 320. Most pixels are labeled one of the eight classes with some unlabeled. All the experiments have been done on a Dual Core, 3GHz Processor with parallel tasks running on both the processor. We have sticked with the simple features used in the state-of-the-art algorithm [10] with replacing the gradient computations by our method in order to demonstrate the success of the proposed algorithm. The experimental method has the following flow chart.

1. We have used a pairwise 4 connected grid model.
2. We found the appropriate order of the Gradient Operator to be 0.7 using the score Function described in Section 2. To run over all images and compute the approximate one decimal approximate of FDOG it took about 5 minutes.
3. For the unary features we first computed RGB intensities of each pixel, along with the normalized vertical and horizontal positions. We expand these initial features into a larger set using sinusoidal expansion, specifically use sin(c.s) and cos(c.s) where s is the initial feature mentioned for all binary vectors c of appropriate lengths leading to total 64 features. [16]
4. Compute FHOG (fractional HOG)i.e. Histogram of Oriented gradient Features but the gradients are found using the proposed method i.e. using a third order approximate kernel using 1 where the order was taken as 0.6.
5. For Edge Features between pixels i and j we considered 21 base features comprsinga constant of one,Euclidian Norm of RGB value difference dicretized into 10 levels and maximum response of FDOG at i or j over three channels and aging binning into 10 dicrete levels and one feature based on the difference of RGB intensities.
6. All methods are trained using TRW to fit the Clique Logistic Loss.

Apart from using FDOG we also change the Optimization algorithm to R-PROP for improved computational efficiency.

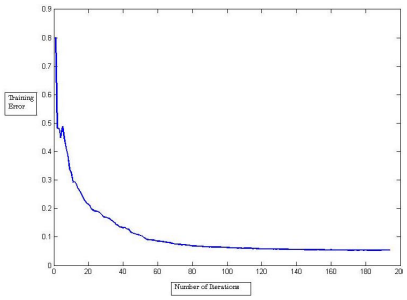

Figure 7. Training Performance of the Model.

Table 2. Comparison with different algorithms

Algorithm	Avg. Accuracy
Region-based energy[16]	65.5
Stacked Labeling [12]	66.2
RGB-D [24]	74.5
TRW BP+Clique Loss [10]	77.9
TRW BP+ CLique Loss+ FDOG [us]	79.2

3.2. Performance

The model is trained with 90 Images in a Monte Carlo Manner with 10 fold cross validation splitting the Training Set into Training and Validation in 90:10 ratio and the data was also shuffled randomly to avoid pathological ordering. The training took around 5 hours on a Intel dual core 3GHz Processor. The training performance is shown in Figure 7

The TRW BP approach by Domke et.al [10] using same set of features but HOG and and edge detectors use DOG is the state of the art 22.1% average error on test data. We bettered that in our experiments reducing the error by almost 10%. Our average error on the test data turned out to be 20.8%. Table 2 shows the result on Stanford Dataset. Some visual results on test data are shown in Figure 8

4. Discussion and Future Work

As can be seen from the demonstrated results we have done quite well on the difficult Stanford Backgrounds Dataset. Though 20.8% error is quite high, there is a lot of room for improvement. The reason for this improvement is the efficient calculation of Gradient Features and it can improve the performance of any of the existing algorithms that are based on Gradient Features. Some of the future work in this might be to extend the concept directly from Frequency Domain and see how it affects the performance. Some, other datasets can be explored. Also, one important option is to apply the Texture Based Approach using this framework as intuitively it seems to be a more realistic idea in case of Contour Based Edge Detection. We believe that using this unified framework the Texton Based approach can also be improved by a few notch.

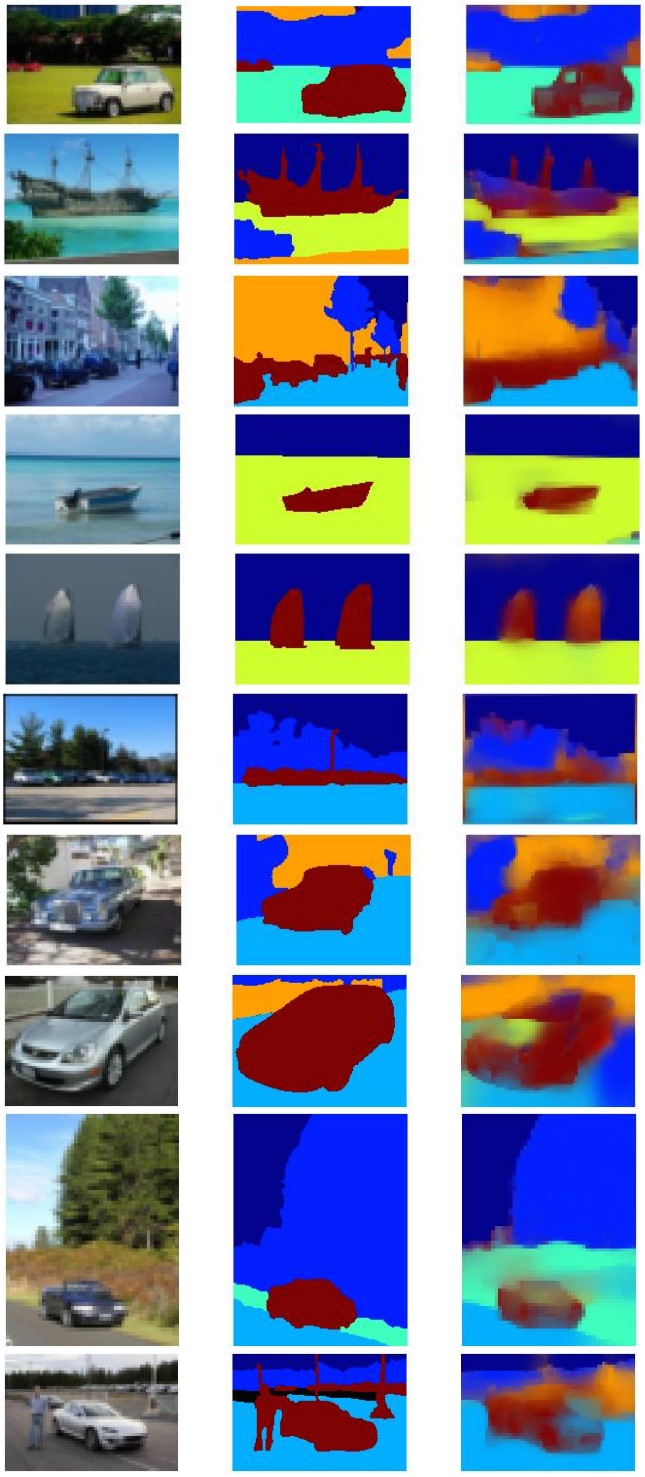

Figure 8. Image segmentation. Each row, from left to right: original image, true label and estimated label.

References

[1] A. Acharya, S. Das, I. Pan, and S. Das. Extending the concept of analog butterworth filter for fractional order systems.

Signal Process., 94:409–420, Jan. 2014. 1, 2

[2] Y. Benjamini and Y. Hochberg. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300, 1995. 3

[3] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, Cambridge, MA, USA, 1987. 1

[4] M. Brown, R. Szeliski, and S. Winder. Multi-image matching using multi-scale oriented patches. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 1 - Volume 01*, CVPR '05, pages 510–517, Washington, DC, USA, 2005. IEEE Computer Society. 2

[5] J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(6):679–698, June 1986. 2, 3

[6] M. A. Carreira-Perpinan and G. E. Hinton. On contrastive divergence learning. Citeseer, 2005. 5

[7] O. Chum and J. Matas. Matching with prosac " progressive sample consensus. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 1 - Volume 01*, CVPR '05, pages 220–226, Washington, DC, USA, 2005. IEEE Computer Society. 2

[8] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893 vol. 1, June 2005. 1

[9] S. Das. Functional fractional calculus, 2nd edition. Springer, 2011. 1

[10] J. Domke. Learning graphical model parameters with approximate marginal inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(10):2454–2467, 2013. 1, 4, 5, 6, 7

[11] S. Gould, R. Fulton, and D. Koller. Decomposing a scene into geometric and semantically consistent regions. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 1–8, Sept 2009. 3

[12] S. Gould, R. Fulton, and D. Koller. Decomposing a scene into geometric and semantically consistent regions. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 1–8. IEEE, 2009. 7

[13] C. Harris and M. Stephens. A combined corner and edge detector. In *Proceedings of the 4th Alvey Vision Conference*, pages 147–151, 1988. 2

[14] S. Kim and I. S. Kweon. Robust model-based scene interpretation by multilayered context information. *Computer Vision and Image Understanding*, 105(3):167–187, 2007. 6

[15] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009. 1, 4

[16] G. Konidaris, S. Osentoski, and P. S. Thomas. Value function approximation in reinforcement learning using the fourier basis. In *AAAI*, 2011. 6, 7

[17] H. S. Koppula, A. Anand, T. Joachims, and A. Saxena. Semantic labeling of 3d point clouds for indoor scenes. In *Advances in Neural Information Processing Systems*, pages 244–252, 2011. 6

[18] E. C. Larson and D. M. Chandler. Most apparent distortion: full-reference image quality assessment and the role of strategy. *Journal of Electronic Imaging*, 19(1):011006–011006, 2010. 3

[19] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, Nov. 2004. 1

[20] J. Malik, S. Belongie, T. Leung, and J. Shi. Contour and texture analysis for image segmentation. *Int. J. Comput. Vision*, 43(1):7–27, June 2001. 1

[21] J. Malik, S. Belongie, J. Shi, and T. Leung. Textons, contours and regions: Cue integration in image segmentation. In *Proceedings of the International Conference on Computer Vision-Volume 2 - Volume 2*, ICCV '99, pages 918–, Washington, DC, USA, 1999. IEEE Computer Society. 1

[22] J. Malik and P. Perona. Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A*, 7:923–932, 1990. 1

[23] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 2, pages 416–423 vol.2, 2001. 3, 4

[24] D. Munoz, J. A. Bagnell, and M. Hebert. Stacked hierarchical labeling. In *Computer Vision–ECCV 2010*, pages 57–70. Springer, 2010. 7

[25] M. Nikolova. Model distortions in bayesian map reconstruction. *Inverse Problems and Imaging*, 1(2):399, 2007. 5

[26] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12(7):629–639, July 1990. 2

[27] B. C. Russell, W. T. Freeman, A. A. Efros, J. Sivic, and A. Zisserman. Using multiple segmentations to discover objects and their extent in image collections. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 1605–1614. IEEE, 2006. 6

[28] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, Aug. 2000. 1

[29] L. Stewart, X. He, and R. S. Zemel. Learning flexible features for conditional random fields. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(8):1415–1426, 2008. 5

[30] J. Tighe and S. Lazebnik. Superparsing: scalable nonparametric image parsing with superpixels. In *Computer Vision–ECCV 2010*, pages 352–365. Springer, 2010. 6

[31] M. J. Wainwright. Estimating the wrong graphical model: Benefits in the computation-limited setting. *The Journal of Machine Learning Research*, 7:1829–1859, 2006. 5

[32] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305, 2008. 6

[33] J. J. Weinman, L. Tran, and C. J. Pal. Efficiently learning random fields for stereo vision with sparse message passing. In *Computer Vision–ECCV 2008*, pages 617–630. Springer, 2008. 5